

On the importance of graph search algorithms for DRGEP-based mechanism reduction methods

Kyle E. Niemeyer^a, Chih-Jen Sung^{b,*}

^a*Department of Mechanical and Aerospace Engineering
Case Western Reserve University, Cleveland, OH 44106, USA*

^b*Department of Mechanical Engineering
University of Connecticut, Storrs, CT 06269, USA*

Abstract

The importance of graph search algorithm choice to the directed relation graph with error propagation (DRGEP) method is studied by comparing basic and modified depth-first search, basic and R -value-based breadth-first search (RBFS), and Dijkstra's algorithm. By using each algorithm with DRGEP to produce skeletal mechanisms from a detailed mechanism for n -heptane with randomly-shuffled species order, it is demonstrated that only Dijkstra's algorithm and RBFS produce results independent of species order. In addition, each algorithm is used with DRGEP to generate skeletal mechanisms for n -heptane covering a comprehensive range of autoignition conditions for pressure, temperature, and equivalence ratio. Dijkstra's algorithm combined with a coefficient scaling approach is demonstrated to produce the most compact skeletal mechanism with a similar performance compared to larger skeletal mechanisms resulting from the other algorithms. The computational efficiency of each algorithm is also compared by applying the DRGEP method with each search algorithm on the large detailed mechanism for n -alkanes covering n -octane to n -hexadecane with 2115 species and 8157 reactions. Dijkstra's algorithm implemented with a binary heap priority queue is demonstrated as the most efficient method, with a CPU cost two orders of magnitude less than the other search algorithms.

Keywords: Mechanism reduction, Skeletal mechanism, DRG, DRGEP, Graph search algorithm, Dijkstra's algorithm

*Corresponding author

Email address: cjsung@engr.uconn.edu (Chih-Jen Sung)

1. Introduction

The directed relation graph (DRG) method for the skeletal reduction of large reaction mechanisms, originally developed by Lu and Law [1, 2], has been shown to be applicable to relevant transportation fuel components [3, 4]. Development of this approach has since focused on variants of the method [5–8], but DRG with error propagation (DRGEP) in particular has received much attention [6, 8–13]. The DRGEP differs mainly from DRG in that it takes error propagation down graph pathways into account. Further improvement of the DRGEP method is motivated by the large and continually increasing size of detailed reaction mechanisms for liquid transportation fuels [14]; for instance, see the recent biodiesel surrogate mechanism of Herbinet et al. [15] that contains 3299 species and 10806 reactions.

The DRG method maps the coupling of species onto a directed graph and finds unimportant species for removal by eliminating weak graph edges using an error threshold, where the graph nodes represent species and graph edge weights indicate the dependence of one species on another. Following a simple graph search initiated at certain preselected target species (e.g., fuel, oxidizer, important pollutants), species unreached by the search are considered unimportant and removed from the skeletal mechanism. DRGEP modifies this approach slightly by considering the dependence of a target on other species due to a certain path as the product of intermediate edge weights and the overall dependence is the maximum of all path dependencies.

One point that has not received much attention is the choice of graph search algorithm used to calculate this overall dependence. While the DRG method only needs to find connected graph nodes and can therefore use any search algorithm, caution must be taken when selecting the method used with DRGEP. The issue of calculating the overall dependence is actually a single-source shortest path problem [16] where the “distance” of the path is the product of intermediate edge weights rather than the sum, and the “shortest” path is that which has the maximum overall dependence rather than the minimum. Search algorithms that in general do not correctly calculate the overall dependence will underestimate the importance of species and cause premature removal from the skeletal mechanism. Further, the resulting skeletal mechanism is independent of the order of the species in the mechanism only if a specific algorithm can correctly capture and calculate

the overall dependence of species. Therefore, the reliability of various algorithms needs to be studied by determining whether each is dependent on the order of species in a detailed mechanism.

The efficiency of the graph search algorithm is also an important factor due to the recent use of DRGEP in dynamic skeletal reduction approaches [10–12]. In the worst case, DRGEP is applied at every spatial grid location and each time step to generate locally relevant skeletal mechanisms, although as Shi et al. [12] demonstrated this can be eased by combining dynamic DRGEP-based reduction with an adaptive multi-grid chemistry model. Regardless, the computational cost of the search algorithm increases with the number of species and must be considered when comparing algorithms due to the large and ever-increasing size of detailed reaction mechanisms.

Most DRGEP studies reported using algorithms based on either depth-first search (DFS) [8] or breadth-first search (BFS) [10–13], but no comparison has been performed. Here we compare such methods with Dijkstra’s algorithm, the classical solution to the single-source shortest path problem [16, 17], in order to demonstrate the weaknesses of DFS- and BFS-based approaches and the subsequent effectiveness and reliability of Dijkstra’s algorithm. First, by randomly shuffling the order of species in the detailed mechanism for n -heptane of Curran et al. [18], we show that DFS- and BFS-based algorithms can generate results dependent on the order of species in the reaction mechanism while Dijkstra’s algorithm generates consistent results regardless of species order. Second, we demonstrate that, for a given error limit, more compact skeletal mechanisms are possible with DRGEP by using Dijkstra’s algorithm. This is done by comparing skeletal mechanisms generated with different search algorithms for n -heptane covering a comprehensive range of autoignition conditions for pressure, temperature, and equivalence ratio. Third, we compare the computational efficiency of the various search algorithms by calculating the CPU time required to perform the graph search on the large detailed mechanism of Westbrook et al. [19] covering oxidation of n -alkanes from n -octane to n -hexadecane containing 2115 species and 8157 reactions.

2. Methodology

2.1. DRGEP Method

In the current work we use the DRGEP method of Pepiot-Desjardins and Pitsch [6], described here in brief. Accurate calculation of the production of a

species A that is strongly dependent on another species B requires the presence of species B in the reaction mechanism. This dependence is expressed with the direct interaction coefficient (DIC):

$$r_{AB} = \frac{\left| \sum_{i=1}^{n_R} \nu_{A,i} \omega_i \delta_B^i \right|}{\max(P_A, C_A)}, \quad (1)$$

where

$$P_A = \sum_{i=1}^{n_R} \max(0, \nu_{A,i} \omega_i), \quad (2)$$

$$C_A = \sum_{i=1}^{n_R} \max(0, -\nu_{A,i} \omega_i), \quad (3)$$

$$\delta_B^i = \begin{cases} 1 & \text{if reaction } i \text{ involves species } B, \\ 0 & \text{otherwise,} \end{cases} \quad (4)$$

A and B represent the species of interest (with dependency in the $A \rightarrow B$ direction meaning that A depends on B), i the i th reaction, $\nu_{A,i}$ the stoichiometric coefficient of species A in the i th reaction, ω_i the overall reaction rate of the i th reaction, and n_R the total number of reactions.

After calculating the DIC for all species pairs, a graph search is performed starting at user-selected target species to find the dependency paths for all species from the targets. A path-dependent interaction coefficient (PIC) represents the error propagation through a certain path and is defined as the product of intermediate DICs between the target T and species B through pathway p :

$$r_{TB,p} = \prod_{j=1}^{n-1} r_{S_j S_{j+1}}, \quad (5)$$

where n is the number of species between T and B in pathway p and S_j is a placeholder for the intermediate species j starting at T and ending at B . An overall interaction coefficient (OIC) is then defined as the maximum of all PICs between the target and each species of interest:

$$R_{TB} = \max_{\text{all paths } p} (r_{TB,p}). \quad (6)$$

Pepiot-Desjardins and Pitsch [6] also proposed a coefficient scaling procedure to better relate the OICs from different points in the reaction system

evolution that we adopt here. A pseudo-production rate of a chemical element a based on the production rates of species containing a is defined as

$$P_a = \sum_{\text{all species } S} N_{a,S} \max(0, P_S - C_S), \quad (7)$$

where $N_{a,S}$ is the number of atoms a in species S and P_S and C_S are the production and consumption rates of species S as given by Eqs. 2 and 3, respectively. The scaling coefficient for element a and target species T at time t is defined as

$$\alpha_{a,T}(t) = \frac{N_{a,T} |P_T - C_T|}{P_a}. \quad (8)$$

For the set of elements $\{\mathcal{E}\}$, the global normalized scaling coefficient for target T at time t is

$$\alpha_T(t) = \max_{a \in \{\mathcal{E}\}} \left(\frac{\alpha_{a,T}(t)}{\max_{\text{all time}} \alpha_{a,T}(t)} \right). \quad (9)$$

Given a set of kinetics data $\{\mathcal{D}\}$ and target species $\{\mathcal{T}\}$, the overall importance of species S to the target species set is

$$\overline{R_S} = \max_{\substack{T \in \{\mathcal{T}\} \\ k \in \{\mathcal{D}\}}} \left[\max_{\text{all time, } k} (\alpha_T R_{TS}) \right]. \quad (10)$$

We employ the same sampling method as given by Niemeyer et al. [8], where constant volume autoignition simulations are performed using SENKIN [20] with CHEMKIN-III [21]. Species are considered unimportant and removed if their $\overline{R_S}$ value falls below a cutoff threshold ε_{EP} , which is selected using an iterative procedure based on a user-defined error limit for ignition delay prediction [8]. Reactions are eliminated if any participating species are removed.

2.2. Graph search algorithms

In this study we compare the results of DRGEP using basic DFS, modified DFS, basic BFS, R -value-based BFS (RBFS), and Dijkstra’s algorithm. Cormen et al. [16] presented detailed discussion of and pseudocode for DFS, BFS, and Dijkstra’s algorithm, while the modified DFS used by Niemeyer et al. [8] and RBFS of Liang et al. [10] differ only slightly from the basic DFS and BFS.

2.2.1. DFS-based algorithms

The DFS initiates at a root node, in this case a target species, and explores the graph edges connecting to other nodes. The first node found is added to a last in, first out stack, then the search moves to this node and repeats. In this manner, the search continues deeper down the graph pathway until it either reaches a node with no connections or all the connecting nodes have been explored, then backtracks one position up the stack. The search is performed separately using each target species as the root node and the maximum OIC is stored for each species. Lu and Law first introduced the DFS in this context for the DRG method [1].

The modified DFS used by Niemeyer et al. [8] for the DRGEP method (but not described in detail) differs from the basic DFS in that the OIC values for all target species are set to unity before starting the search at the first target. The search then only repeats starting at other target species not discovered in the initial search. The resulting OIC values combine the dependence of all targets on the species and this prevents the use of coefficient scaling based on individual target species activity.

2.2.2. BFS-based algorithms

The BFS initiates at the root node and explores all adjacent graph edges, adding connected nodes to a first in, first out queue. After discovering all connected nodes, the search moves to the first node in the queue and restarts the procedure, moving to the next node in the queue after discovering all connected nodes. Previously discovered nodes are not repeated. As with the DFS, the search initializes at each target separately and the maximum OIC for each species is stored.

The R -value-based BFS (RBFS), first introduced by Liang et al. [10, 11] and also used by Shi et al. [12, 13], differs from the standard BFS in that PICs smaller than the error threshold ε_{EP} are not explored. In other words, the search ignores graph pathways that would lead to an OIC below the cut-off value and only allows discovery of important pathways. This increases efficiency by avoiding exploration of unimportant pathways but causes the RBFS to depend on the value of ε_{EP} , unlike the other search methods. As such, the primary application of this algorithm lies in dynamic reduction where the threshold value is known a priori, rather than general comprehensive skeletal reduction where the threshold is to be determined iteratively based on a user-defined error limit [8].

2.2.3. Dijkstra’s algorithm

Dijkstra’s algorithm, originally introduced by Dijkstra [17] and discussed in detail by Cormen et al. [16], differs from DFS, BFS, and their variants because it was specifically designed to calculate the shortest paths from a single source node to all other nodes rather than simply search the graph to find connected nodes. As mentioned previously, calculating OIC values is a shortest path problem where the “shortest” path is that with the maximum product of intermediate edge weights. This is the only modification needed to apply Dijkstra’s algorithm to the current problem. In brief, Dijkstra’s algorithm functions similarly to BFS except that it starts with the set of graph nodes stored in a max-priority queue. The algorithm finds and removes the node with the maximum OIC value (initially, the root node) from the queue and explores adjacent nodes to calculate or update OIC values. When all neighboring nodes are explored, the procedure restarts at the node with the next-highest OIC in the queue. The search completes when the queue is empty. As with previously-discussed search algorithms, this is performed separately for each target species as root node. See the [Appendix](#) for pseudocode describing the algorithm as applied to the DRGEP method.

Transportation and internet routing applications routinely use Dijkstra’s algorithm and so much effort has been focused on optimizing its performance. A naive implementation requires a running time of $O(V^2)$ where V is the number of nodes. For sparse graphs, i.e. where the number of edges E follows $E \sim O(V^2/\log V)$, the runtime can be improved significantly to $O(E \log V)$ by using an adjacency list to store the immediate neighbors of nodes and a binary max-heap priority queue [16]. Using a Fibonacci heap can reduce the runtime further to $O(E + V \log V)$ [16, 22], although in many cases binary heaps prove more efficient [23]. These will be compared in order to determine the most efficient implementation for DRGEP. Briefly, a binary max-heap is a tree where each node has at most two child nodes and the associated key (in this case, OIC) is always greater than or equal to the keys of the child nodes. By storing the nodes in this manner, a costly search is not required to find the node with the highest key during the operation of Dijkstra’s algorithm. Fibonacci heaps function similarly, but rather than a single tree, the heaps are made up of a collection of trees with more relaxed structures designed to avoid unnecessary operations. Detailed discussion of the heap implementations is beyond the scope of this paper and we refer interested readers to Cormen et al. [16].

3. Results and discussion

3.1. Reliability of graph search algorithms

In order to demonstrate the dependence of DFS- and BFS-based methods on the order of species in a reaction mechanism, all five graph search algorithms are used with DRGEP to generate skeletal mechanisms from the detailed mechanism for *n*-heptane of Curran et al. [18], containing 561 species and 2539 reactions, where the order of species is randomly shuffled. Basic DFS, basic BFS, RBFS, and Dijkstra’s algorithm are used both with and without the target-based coefficient scaling while the modified DFS is not compatible with this and so is used without it. Autoignition chemical kinetics data are sampled from initial conditions at 1000 K, 1 atm, and equivalence ratios of 0.5–1.5. Oxygen, nitrogen, *n*-heptane, and the hydrogen radical are selected as target species.

Figure 1 shows the ignition delay error of skeletal mechanisms at varying levels of detail generated by DRGEP using RBFS and Dijkstra’s algorithm with coefficient scaling and the modified DFS. Basic DFS and BFS demonstrate similar dependence on species order to the modified DFS and thus are omitted from Fig. 1 for clarity, but the skeletal mechanism sizes from each may be different as will be shown. Similarly, the coefficient scaling comparison is also omitted because its application does not affect dependence on species order, although the resulting skeletal mechanism sizes are different. Comparison of results from the original mechanism and five mechanisms with randomly shuffled species illustrates the dependence of the modified DFS on the order of species in the mechanism while RBFS and Dijkstra’s algorithm produce consistent results regardless of species order. This is because DFS and BFS explore graphs based on the order of nodes while Dijkstra’s algorithm follows the strongest path independent of order [16]. In addition, Fig. 1 shows that both Dijkstra’s algorithm and RBFS produce smaller skeletal mechanisms before the error becomes unacceptably large.

The dependence of DFS, modified DFS, and BFS on species order stems from the fact that these algorithms do not in general calculate the correct OIC for every species and as such can underestimate the importance of species, causing unwarranted removal. RBFS produces the same results as Dijkstra’s algorithm because it prevents exploration of unimportant paths, i.e. paths with PICs (see Eq. 5) less than the threshold, and therefore only discovers paths that lead to an OIC greater than the threshold value. While this value may be smaller than the correct OIC as calculated by Dijkstra’s algorithm,

the species is nonetheless considered important and retained in the skeletal mechanism. This subtle error could be significant, though, if the DRGEP method is followed by a further reduction stage such as sensitivity analysis that depends on the OIC values [8].

3.2. Effectiveness of graph search algorithms

Skeletal mechanisms for *n*-heptane are generated covering a comprehensive range of conditions using basic DFS, modified DFS, basic BFS, and Dijkstra’s algorithm. Basic DFS, BFS, and Dijkstra’s algorithm are used with and without the coefficient scaling in order to determine its effect as well. Autoignition chemical kinetics data are sampled from initial conditions at 600–1600 K, 1–20 atm, and equivalence ratios of 0.5–1.5. Oxygen, nitrogen, *n*-heptane, and the hydrogen radical are selected as target species, and the error limit in ignition delay prediction is 30%. RBFS is not compared here because it is not suited for a comprehensive skeletal reduction due to its dependence on threshold value; additionally, based on the results in the previous section we assume the resulting skeletal mechanism would match that from Dijkstra’s algorithm.

The results using the original mechanism of Curran et al. [18] are summarized in Table 1. All search methods lead to skeletal mechanisms with similar performance, but Dijkstra’s algorithm with coefficient scaling produces the most compact skeletal mechanism. BFS and modified DFS produce similar results while basic DFS is unable to generate a comparable skeletal mechanism for the given error limit. This weakness is most likely due to the fact that basic DFS finds only the first path from target to species, which typically contains many intermediate species and severely underestimates the OICs for some important species. The modified DFS performs better by inserting the other targets into the pathways and therefore increasing the OIC values for important species, while BFS finds the path with the shortest number of intermediate nodes [16].

Table 1 also demonstrates the need for coefficient scaling. Both basic DFS and Dijkstra’s algorithm generate more compact skeletal mechanisms with the scaling, while BFS actually produces a slightly larger skeletal mechanism. Without scaling, Dijkstra’s algorithm is unable to produce a more compact skeletal mechanism than the other methods, despite its ability to generate results independent of species order in the parent mechanism.

3.3. Efficiency of graph search algorithms

The computational costs of DFS, modified DFS, basic BFS, RBFS, and Dijkstra’s algorithm are compared by applying the DRGEP method to the detailed mechanism of Westbrook et al. [19] for n -alkanes covering n -octane through n -hexadecane, which contains 2115 species and 8157 reactions. Kinetics data are sampled from constant volume autoignition of n -decane with initial conditions covering 600–1600 K, 1–20 atm, and equivalence ratios of 0.5–1.5. The efficiency improvements available to Dijkstra’s algorithm including use of adjacency list, binary heap, and Fibonacci heap are also compared to the naive implementation.

The average computational costs of DFS, modified DFS, BFS, Dijkstra’s algorithm, and its improvements are listed in Table 2. In order to perform a fair comparison, a version of Dijkstra’s algorithm modified to similarly depend on the threshold value is compared with RBFS. Figure 2 shows the average computational cost of RBFS, Dijkstra’s algorithm, and its improvements as a function of threshold value. The most important result shown in Table 2 and Fig. 2 is that Dijkstra’s algorithm implemented with binary heaps is faster than all other search algorithms, and is in fact more efficient than the same algorithm implemented with Fibonacci heaps even though the theoretical time limit with binary heaps is higher. This result is consistent with literature comparisons for sparse graphs [24, 25].

4. Concluding remarks

Graph search algorithms used in the DRGEP method for skeletal mechanism reduction are compared. Basic DFS, basic BFS, modified DFS, RBFS, and Dijkstra’s algorithm are implemented in DRGEP and used to generate (1) skeletal mechanisms covering a limited range of conditions using a randomly shuffled detailed mechanism for n -heptane to determine the dependence of results on species order and (2) skeletal mechanisms covering a comprehensive range of conditions to determine the effectiveness of each algorithm. The RBFS algorithm is not used to generate a comprehensive skeletal mechanism because it is not suitable for use when the cutoff threshold is not known a priori. Both Dijkstra’s algorithm and RBFS are able to generate consistent results independent of species order, and Dijkstra’s algorithm used with target-based coefficient scaling generates a more compact skeletal mechanism than the other methods. Even with the improved search algorithm, however, the size of the resulting skeletal mechanism (131 species and 651

reactions) is larger than that of DRGEP with sensitivity analysis (108 species and 406 reactions) [8], suggesting that the post-DRGEP sensitivity analysis is still required.

In addition to the reliability and effectiveness of the graph search algorithms, the efficiency is also compared due to its importance to dynamic DRGEP approaches. Efficiency improvements available to Dijkstra’s algorithm are also implemented and compared to the other search algorithms. Dijkstra’s algorithm with a binary heap priority queue runs two orders of magnitude faster than the other search methods and is the most efficient implementation of Dijkstra’s algorithm. As such, this approach is recommended for use in dynamic skeletal reduction using DRGEP. Dynamic approaches that avoid entirely repeating the search by updating graph edge values have also been developed [23, 26] and will be the focus of future work.

Acknowledgments

This work was supported by the National Science Foundation under Grant No. 0932559 and the Department of Defense through the National Defense Science and Engineering Graduate Fellowship program.

Appendix

The following pseudocode describes Dijkstra’s algorithm for use in the DRGEP method to calculate the OICs for all species relative to a target species T using the set of DIC values r_{AB} , adapted from Cormen et al. [16]. The adjacency list adj contains the list of adjacent (directly connected) species in the directed graph. Operations are left in general form to be applicable to other implementations of Dijkstra’s algorithm using heaps as well as the basic version. The procedure MAKEQ (Q) generates the max-priority queue Q , MAXQ (R , Q) returns the node with the highest OIC value, REMQ (u , Q) removes the node u from Q , and UPDATE (R , v , R_{tmp}) increases the OIC value for node v to R_{tmp} .

subroutine Dijkstra (r_{AB} , adj , T)

```

 $R \leftarrow$  zeros
 $R(T) \leftarrow 1.0$ 
call MAKEQ ( $Q$ )
while  $Q \neq \emptyset$  do
     $u \leftarrow$  MAXQ ( $R$ ,  $Q$ )

```

```

call REMQ ( $u$ ,  $Q$ )
for each node  $v \in adj(u)$  do
     $R_{tmp} \leftarrow R(u) \cdot r_{uv}$ 
    if  $R_{tmp} > R(v)$  then
        call UPDATE ( $R$ ,  $v$ ,  $R_{tmp}$ )
    end if
end for
end while
return  $R$ 
end subroutine Dijkstra

```

References

- [1] T. Lu, C. K. Law, Proc. Combust. Inst. 30 (2005) 1333–1341.
- [2] T. Lu, C. K. Law, Combust. Flame 146 (2006) 472–483.
- [3] T. Lu, C. K. Law, Combust. Flame 144 (2006) 24–36.
- [4] T. Lu, C. K. Law, Combust. Flame 154 (2008) 153–163.
- [5] X. L. Zheng, T. Lu, C. K. Law, Proc. Combust. Inst. 31 (2007) 367–375.
- [6] P. Pepiot-Desjardins, H. Pitsch, Combust. Flame 154 (2008) 67–81.
- [7] W. Sun, Z. Chen, X. Gou, Y. Ju, Combust. Flame 157 (2010) 1298–1307.
- [8] K. E. Niemeyer, C. J. Sung, M. P. Raju, Combust. Flame 157 (2010) 1760–1770.
- [9] I. G. Zsély, T. Nagy, J. M. Simmie, H. J. Curran, in: 4th European Combustion Meeting, 810045, 2009.
- [10] L. Liang, J. Stevens, J. T. Farrell, Proc. Combust. Inst. 32 (2009) 527–534.
- [11] L. Liang, J. G. Stevens, S. Raman, J. T. Farrell, Combust. Flame 156 (2009) 1493–1502.
- [12] Y. Shi, L. Liang, H.-W. Ge, R. D. Reitz, Combust. Theor. Model. 14 (2010) 69–89.

- [13] Y. Shi, H.-W. Ge, J. L. Brakora, R. D. Reitz, *Energy Fuels* 24 (2010) 1646–1654.
- [14] T. Lu, C. K. Law, *Prog. Energy Comb. Sci.* 35 (2009) 192–215.
- [15] O. Herbinet, W. J. Pitz, C. K. Westbrook, *Combust. Flame* 157 (2010) 893–908.
- [16] T. H. Cormen, C. E. Leiserson, R. L. Rivest, C. Stein, *Introduction to Algorithms*, MIT Press, Cambridge, MA, 2nd edition, 2001.
- [17] E. W. Dijkstra, *Numer. Math.* 1 (1959) 269–271.
- [18] H. Curran, P. Gaffuri, W. Pitz, C. K. Westbrook, *Combust. Flame* 114 (1998) 149–177.
- [19] C. K. Westbrook, W. J. Pitz, O. Herbinet, H. J. Curran, E. J. Silke, *Combust. Flame* 156 (2009) 181–199.
- [20] A. E. Lutz, R. J. Kee, J. A. Miller, SENKIN: A FORTRAN program for predicting homogeneous gas phase chemical kinetics with sensitivity analysis, Sandia National Laboratories Report No. SAND 87-8248, 1988.
- [21] R. J. Kee, F. M. Rupley, E. Meeks, J. A. Miller, CHEMKIN-III: A FORTRAN chemical kinetics package for the analysis of gas-phase chemical and plasma kinetics, Sandia National Laboratories Report No. SAND 96-8216, 1996.
- [22] M. L. Fredman, R. E. Tarjan, *J. ACM* 34 (1987) 596–615.
- [23] D. Wagner, T. Willhalm, *Lect. Notes Comput. Sci.* 4393 (2007) 23–36.
- [24] B. Cherkassky, A. Goldberg, T. Radzik, *Math. Program.* 73 (1996) 129–174.
- [25] A. V. Goldberg, R. E. Tarjan, Expected performance of Dijkstra’s shortest path algorithm, Technical Report TR-530-96, NEC Research Institute, Princeton University, 1996.
- [26] L. Roditty, U. Zwick, *Lect. Notes Comput. Sci.* 3221 (2004) 580–591.

Algorithm		# Species	# Reactions	Max. Error
DFS	no scaling	461	2304	19%
	scaling	449	2267	19%
mod. DFS		173	868	28%
BFS	no scaling	180	891	29%
	scaling	207	921	25%
Dijkstra	no scaling	178	986	23%
	scaling	131	651	17%

Table 1: Comparison of *n*-heptane skeletal mechanism sizes generated by the DRGEP method using DFS, BFS, and Dijkstra’s algorithm with and without coefficient scaling and modified DFS. The original mechanism contains 561 species and 2539 reactions [18].

Algorithm	Mean cost (sec)	Cost normalized by Dijkstra
DFS	1.48×10^{-1}	2.20
mod. DFS	1.46×10^{-1}	2.18
BFS	8.30×10^{-2}	1.24
Dijkstra (naive)	6.72×10^{-2}	1.00
Dijkstra (adj)	1.18×10^{-2}	0.176
Dijkstra (binary heap)	1.5×10^{-3}	0.028
Dijkstra (Fibonacci heap)	3.4×10^{-3}	0.051

Table 2: Comparison of average CPU time costs using DFS, modified DFS, BFS, and Dijkstra’s algorithm, which includes the naive implementation, use of adjacency list, binary heaps, and Fibonacci heaps. The kinetics data used are generated from *n*-decane autoignition over a range of initial temperatures and pressures, and at varying equivalence ratios, from a detailed mechanism consisting of 2115 species and 8157 reactions [19].

List of Figures

- 1 Error in autoignition delay prediction as a function of number of species for skeletal mechanisms of *n*-heptane generated by DRGEP using Dijkstra’s algorithm and RBFS with coefficient scaling and modified DFS. The solid lines indicate the results using the original detailed mechanism while the dashed lines indicate the results obtained from five versions with randomly-shuffled species. All 12 results using Dijkstra’s algorithm and RBFS coincide. 16
- 2 Comparison of average CPU time cost as a function of threshold value using RBFS and Dijkstra’s algorithm, which includes the naive implementation, use of adjacency list, binary heaps, and Fibonacci heaps. The kinetics data used are generated from *n*-decane autoignition over a range of initial temperatures and pressures, and at varying equivalence ratios, from a detailed mechanism consisting of 2115 species and 8157 reactions [19]. 17




Figure 1: Error in autoignition delay prediction as a function of number of species for skeletal mechanisms of *n*-heptane generated by DRGEP using Dijkstra's algorithm and RBFS with coefficient scaling and modified DFS. The solid lines indicate the results using the original detailed mechanism while the dashed lines indicate the results obtained from five versions with randomly-shuffled species. All 12 results using Dijkstra's algorithm and RBFS coincide.




Figure 2: Comparison of average CPU time cost as a function of threshold value using RBFS and Dijkstra’s algorithm, which includes the naive implementation, use of adjacency list, binary heaps, and Fibonacci heaps. The kinetics data used are generated from *n*-decane autoignition over a range of initial temperatures and pressures, and at varying equivalence ratios, from a detailed mechanism consisting of 2115 species and 8157 reactions [19].